

CLAIM AMENDMENTS

This listing of claims will replace all prior versions, and listings, of claims in the application.

1 1. (Currently Amended) A method of reading data comprising the steps of:
2 receiving a request for a stripe of erasure coded data stored across a
3 plurality of storage devices, the stripe comprising stripe blocks;
4 sending read messages to at least a quorum of the storage devices;
5 receiving at least the quorum of reply messages from the storage
6 devices, the quorum of the reply messages ~~including at least~~being greater than
7 a minimum number of the stripe blocks needed to decode the stripe of erasure
8 coded data, and the quorum meeting a quorum condition of a number such that
9 any two selections of the number of the stripe blocks intersect in the minimum
10 number of the stripe blocks; and
11 in response to receiving at least the quorum of reply messages and not
12 necessarily a reply from each of the storage devices of the plurality, decoding
13 the stripe of erasure coded data from at least the minimum number of the
14 stripe blocks, thereby forming the data.

1 2. (Currently Amended) The method of claim 1 wherein each of the reply
2 messages is included within the quorum only if it indicates that there is no
3 pending write for the stripe block stored on the storage device associated with
4 the reply message.

1 3. (Currently Amended) The method of claim 1 wherein each of the reply
2 messages is included within the quorum only if it indicates that the stripe
3 block associated with the reply message has a timestamp that matches other
4 timestamps associated with other reply messages within the quorum.

1 4. (Currently Amended) The method of claim 1 wherein the storage devices
2 are distributed with each storage device comprising a CPU and storage and
3 wherein a coordinator performs the steps.

1 5. (Original) The method of claim 4 wherein the coordinator comprises one
2 of the storage devices.

1 6. (Original) The method of claim 5 wherein the coordinator effectively
2 sends one of the read messages to itself.

1 7. (Original) The method of claim 6 wherein the coordinator effectively
2 receives one of the reply messages from itself.

1 8. (Currently Amended) A method of reading data comprising the steps of:
2 receiving a request for a stripe of erasure coded data stored across a
3 plurality of storage devices, the stripe comprising stripe blocks which
4 comprise a first number of data blocks and a second number of parity blocks;
5 sending read messages to the storage devices;
6 receiving at least a quorum of reply messages from the storage devices
7 which indicate that there is no pending write for the stripe block stored on the
8 storage device, ~~the quorum of the reply messages including at least the first~~
9 ~~number of the stripe blocks~~, the quorum comprising at least the first number
10 plus a half of the second number and the quorum being greater than a
11 minimum number of the stripe blocks needed to decode the stripe of erasure
12 coded data; and
13 in response to receiving at least the quorum of reply messages and not
14 necessarily a reply from each of the storage devices of the plurality, decoding
15 the stripe of erasure coded data from the first number of the stripe blocks,
16 thereby forming the first number of the data blocks.

1 9. (Original) The method of claim 8 wherein the quorum of the reply
2 messages includes validation timestamps which match.

1 10. (Currently Amended) The method of claim 8 wherein the storage devices
2 are distributed with each storage device comprising a CPU and storage and
3 wherein a coordinator performs the steps.

- 1 11. (Original) The method of claim 10 wherein the coordinator comprises
2 one of the storage devices.
- 1 12. (Original) The method of claim 11 wherein the coordinator effectively
2 sends one of the read messages to itself.
- 1 13. (Original) The method of claim 12 wherein the coordinator effectively
2 receives one of the reply messages from itself.
- 1 14. (Currently Amended) The method of claim ~~[[8]]~~10 wherein the
2 coordinator is not one of the storage devices upon which the stripe of erasure
3 coded data is stored.
- 1 15. (Currently Amended) The method of claim 8 further comprising the step
2 of identifying a group of the storage devices as targets from among the
3 plurality of storage devices.
- 1 16. (Original) The method of claim 15 wherein the step of identifying the
2 targets randomly picks the targets.
- 1 17. (Original) The method of claim 15 wherein each of the query messages
2 sent to the targets identifies the storage device as one of the targets.
- 1 18. (Original) The method of claim 17 wherein the reply messages from the
2 targets include the stripe blocks.
- 1 19. (Original) The method of claim 8 wherein the storage devices comprise a
2 distributed storage system.
- 1 20. (Original) The method of claim 19 wherein the distributed storage system
2 comprises a quantity of the storage devices.

1 21. (Original) The method of claim 20 wherein the quantity of the storage
2 devices corresponds to the first number of the data blocks plus the second
3 number of the parity blocks.

1 22. (Original) The method of claim 20 wherein the quantity of the storage
2 devices exceeds the first number of the data blocks plus the second number of
3 the parity blocks.

1 23. (Original) The method of claim 22 further comprising the step of
2 identifying the storage devices upon which the stripe of erasure coded data is
3 stored.

1 24. (Original) The method of claim 8 wherein each of the storage devices
2 comprises a log, wherein the log comprises log entries of each successful write
3 of data, the log entries comprising a stripe indicator, a write timestamp, and a
4 physical location of the stripe block on the storage device.

1 25. (Original) The method of claim 8 wherein the read messages include a
2 stripe indicator.

1 26. (Original) The method of claim 8 wherein the reply messages include a
2 stripe indicator.

27-28. (Canceled)

1 29. (Currently Amended) A computer readable memory comprising
2 computer code for implementing a method of reading a stripe of erasure coded
3 data, the method of reading the stripe of erasure coded data comprising the
4 steps of:

5 receiving a request for a stripe of erasure coded data stored across a
6 plurality of storage devices, the stripe comprising stripe blocks;
7 sending read messages to at least a quorum of the storage devices;
8 receiving at least the quorum of reply messages from the storage
9 devices, the quorum of the reply messages including at least being greater than

10 a minimum number of the stripe blocks needed to decode the stripe of erasure
11 coded data, and the quorum meeting a quorum condition of a number such that
12 any two selections of the number of the stripe blocks intersect in the minimum
13 number of the stripe blocks; and
14 in response to receiving at least the quorum of reply messages and not
15 necessarily a reply from each of the storage devices of the plurality, decoding
16 the stripe of erasure coded data from at least the minimum number of the
17 stripe blocks, thereby forming the data.

1 30. (Original) The computer readable memory of claim 29 wherein each of
2 the reply messages within the quorum indicate that there is no pending write
3 for the stripe block stored on the storage device associated with the reply
4 message.

1 31. (Original) The computer readable memory of claim 29 wherein each of
2 the reply messages within the quorum indicate that the stripe block associated
3 with the reply message has a timestamp that matches other timestamps
4 associated with other reply messages within the quorum.

32. (Canceled)

1 33. (New) The method according to claim 1, wherein in response to receiving
2 fewer than the quorum of replies, the stripe of erasure coded data is recovered
3 by writing the stripe of erasure coded data to a number of the storage devices
4 that at least meets the quorum condition.

1 34. (New) The method according to claim 8, wherein in response to receiving
2 fewer than the quorum of replies, the stripe of erasure coded data is recovered
3 by writing the stripe of erasure coded data to a number of the storage devices
4 that at least meets the quorum condition.